



Should Seed Investors Follow On?

Abraham Othman Ph.D. — Head of Data Science, AngelList

September 1, 2019

Many seed investors have the opportunity to participate in the follow-on equity rounds that a company raises. Using AngelList's large database of seed-round investments, we show that always making those follow-on investments has a higher mean (i.e., expected return), while ignoring follow-on rounds and instead making more seed-stage investments has a higher median (i.e., typical return). We also show that a strategy that selectively filters for when to participate in follow-on rounds based on the seed investment's increase in valuation ("doubling down on winners") does not demonstrably outperform either extreme, although in all cases the performance of the strategies is close and appears to mostly be governed by randomness.

Introduction

Seed-stage investments are typically made using a convertible instrument (either a convertible note or, increasingly, a SAFE)¹. Startups prefer to raise these instruments as their first source of capital because they are relatively straightforward and speedy to issue. At a company's priced equity round, typically a Series A round led by a VC firm, these convertible investments become shares of the company at a discounted valuation, rewarding the seed investors for their investment duration and foresight.

¹ Sometimes startups will raise their first round of capital in a priced investment round, typically called a "Series Seed". However, it is also the case that a "Series Seed" priced round can follow one or more rounds of convertible issuance. Since in general we cannot tell if a "Series Seed" was predated by convertible instruments, we focus our analysis here only on startups raising convertibles before their first priced equity round.



When this conversion event happens, seed-stage investors will often have the opportunity to “follow on” their investment by writing an additional check into the Series A round². *Should seed-stage investors take advantage of these follow-on rights?*

Follow-On Investment Strategies

We use past AngelList data to simulate investor performance from either following on their converting investments, or instead using their capital to invest in the seed round of a new startup. We will look at three different ways that seed-stage investors can deal with the possibility of making follow-on investments:

Never follow on The canonical “spray and pray” strategy where an investor never makes follow-on investments, and instead uses that capital to make additional seed-stage investments.

Always follow on Make follow-on investments whenever a seed-stage investment converts for a positive return³.

Double down on winners Follow on only when the value of your initial investment at least doubles on conversion.

For brevity, we will dub these “Never”, “Always”, and “Double”. Never makes the most unique investments with the fewest amount of money behind each of them. Always makes the fewest unique investments with the most money behind each of them. Double is in the middle between these two extremes.

Three Possible Worldviews

Here are the “stories” of three different generative world-models. Each world-model suggests a different strategy is likely to do well:

² These follow-on rights can be granted informally through their relationship with the company’s founders or other investors, or formally through an investor rights agreement or side letter.

³ Because convertible instruments are almost always priced at a discount to the next priced round, cases in which conversions happen at a loss but the company is a viable and ongoing concern are rare. Exploratory data analysis suggested that investment mark-downs in conversion events were typically part of an intermediate step in a company wind down; it is likely that investors would have neither the inclination nor the ability to invest in these replacement instruments and so we have ignored them in this analysis.



In a **Mean-Reverting** world, hot deals get bid up without any respect to underlying company quality. In this case, marking a high return at Series A is probably a poor future indicator of performance and some of the worst investments you can make (on a forward basis) are in investments that have recently done quite well. In a mean-reverting world, Never should outperform Always, which should significantly outperform Double.

In a **Totally Random** world, past performance has no bearing on future results. Investment performance on a forward basis should be another draw from the heavy-tailed distribution of early-stage investments. In a Totally Random world, we would expect all three strategies to have about the same level of mean performance, but because Never makes more unique investments, it should have a higher median and win the majority of head-to-head comparisons against Always and Double.

In a **Momentum** world, past performance is a positive indicator of future results, and so putting more money into the (currently) best performing investments will be a winning strategy. In this case, we should expect Double to outperform Always, which should outperform Never.

By simulating the performance of the three strategies, we can gather examine the consistency of the actual early-stage investment world with each of these three stylized world models.

Simulating the Strategies

We are interested in seed (and pre-seed) investments, so we filter down AngelList's extensive database of early-stage investments to only convertible notes or SAFEs that are either uncapped, or that have a conversion price of less than \$20 million dollars. Since, especially when uncapped, these instruments can also be used for "bridge rounds" between priced equity rounds we additionally exclude all the convertible instruments created after a company has had a priced equity round. Finally, for the purpose of simulation we restrict our universe to only those investments made in 2014, 2015, 2016, or 2017. (This means that the youngest investments we consider have had at least 18 months to season.) All told, our filtering leaves us with a universe of 1,218 seed-stage investments to consider.

Our simulation proceeds as follows. We start by fixing the number of investments an investor wants to make each year, say ten. Then in each year (2014, 2015, 2016, and 2017) of simulation, the investor draws a list of ten potential investments.



We simulate Never by having the investor write equal-sized checks into all of these potential seed-stage investments. But we cannot necessarily have Always or Double make all of the potential investments because some of their past investments might convert, reducing the investor's budget for new investments. In this case, we randomly choose potential seed-stage investments for the strategy to omit.

For simplicity we assume that investors follow-on with an identical-size check to the one that they wrote originally into the seed round. In practice, the actual amount an investor may have opportunity to invest may be significantly different, either smaller or larger, than their initial investment, depending on, among other factors, the number of different convertible instruments a company raises before their first priced equity round, the number of shares sold at Series A, and if there are any changes in the size of the employee options pool associated with the fundraising. Moreover, the investor may not have explicit legal rights to follow on to their investment and their stake may be based on negotiations with the company founders or other investors at conversion time that result in smaller or larger offered stakes. Our choice to model follow-on investments equal in size to the original seed-round investment reflects this considerable uncertainty⁴.

Even though all the three strategies, Never, Always, and Double, start from the same universe of potential investments, only the Never strategy will make all of them, while Always and Double will make a subset of the investments. This simulation methodology means that, for a specific simulation and investment universe, the observed differences in performance will be directly caused by the merits of the individual strategies. Consequently, we can consider the direct, pairwise rankings between strategies in a single simulation by looking, in retrospect, which strategy would have best served investors seeing that random slice of possible seed investments.

Results

We ran 10,000 simulations of the performance of the strategies targeting ten investments per year. We experimented with different numbers of target investments; the biggest change from increasing the number of investments is an increase in median strategy performance, which we is consistent with what we would expect from the heavy-tailed power-law distribution of investment returns.

In both simulations, Never has the highest median and wins most of the head-to-head simulations (fixing the same universe of potential investments) against both Always and Double. Always has the highest mean.

⁴ We also analysed the effect of having follow-on checks be proportionate to the increase in valuation instead of equally sized to the original checks; this did not qualitatively change our results.



Table 1. Performance of each of the three follow-on investment strategies when making ten investments per year over 10,000 simulations.

Strategy	Mean TVPI	Median TVPI
Never	1.81x	1.60x
Always	1.90x	1.54x
Double	1.80x	1.59x

Results shown in this table are gross values and do not reflect the impact of fees or carry that would reduce returns achievable by an investor.

Table 2. Fraction of simulations where each follow-on investment strategy outperforms the others. Each comparison is made head-to-head and involves the same potential investment universe.

Never	54%	46%	Always
Never	43%	41%	Double
Double	54%	46%	Always

Never and Double had identical results in 16% of simulations.

Double outperforms Always head-to-head and by median. The magnitude of the divergence between Never and Double against Always in the median simulation is smaller at ten investments than at three.

The following scatter plot shows the simulated performance of Always and Never making ten investments per year. The diagonal line divides the plot into two triangles, points in the triangle at upper left represent simulations where Never outperforms Always, while points in the triangle at the lower right represent simulations where Always outperforms Never.

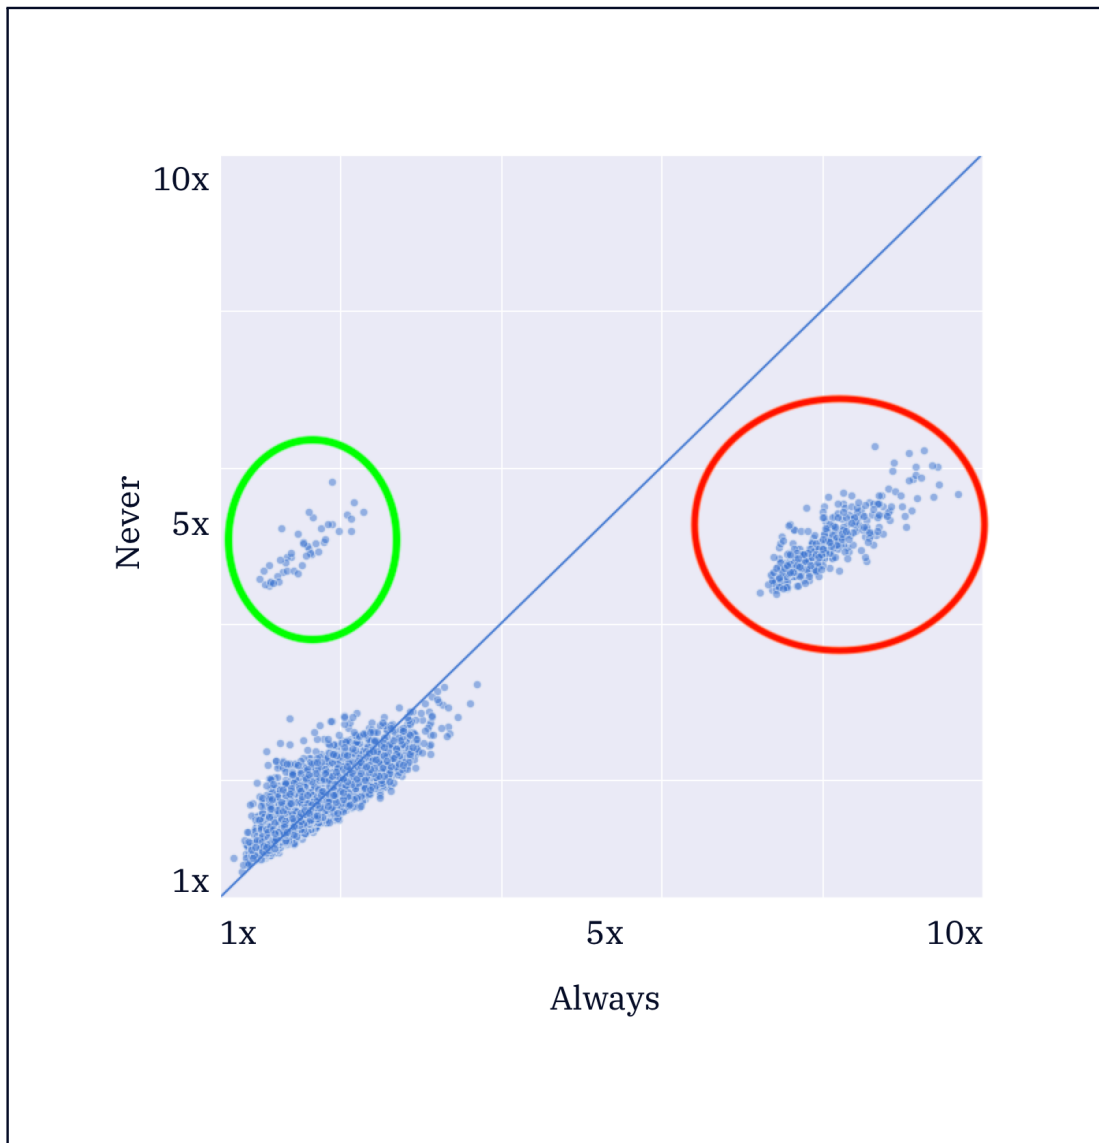


Fig. 1. Scatter plot comparing the gross TVPI of the Always (x axis) and Never (y axis) follow-on strategies. Each dot represents the result of a simulation using identical investment universes. The gross TVPI values in this plot do not reflect the impact of fees or carry that would reduce returns achievable by an investor.

Somewhat surprisingly, the plot cleaves neatly into three clusters, and following the power law of early-stage returns, those three clusters are derived from the simulation and strategy's relationship to the best-performing company on the AngelList platform.



The large cluster at the lower left represents the typical performance of the two strategies; the points in that cluster come from the bulk of simulations that never consider an investment in the seed stage of AngelList's best-performing investment. Observe that in that cluster, Always has returns that have higher variance, while Never has more well-behaved performance. This is consistent with the power law of early-stage investment returns rewarding making more, putatively independent, investments.

The smaller cluster in the red circle at right are simulations where Never performs well but Always performs even better. This cluster is entirely due to Always putting additional capital into the Series A of the best-performing company on the AngelList platform, which company raised a convertible instrument that converted at a modest, roughly 10% return to a Series A before going on to a better than 100x unrealized return after its first priced equity round,

Finally, the smallest cluster, in the green circle, are those simulations where Always has typical performance but Never does well. These are the simulations where following on to a prior investment precludes Always from making the original investment into the best-performing company on the platform. This green circle is a graphical depiction of one of our key findings: that there exists an *opportunity cost* of doing follow-on investments. This is closely related to the power law of early-stage investment returns, which as a general rule suggests that having more winning early-stage investments is always better from the perspective of risk-adjusted returns.

Our experimental results are ambiguous as to whether Double's outperformance of Always in the head-to-head and median measures is because it makes more investments (like Never) or because its selection of follow-on investments is superior to Always. To study this issue we consider those follow-on investments with valuation changes in the table below:

Compared to the universe of all follow-on investments, the Series A investments that followed the doubling of a convertible have a higher median, lower mean, and identical fraction of "Went to Zero" investments. Investor preferences for mean or median return would seem to support either of these strategies, and it does not seem reasonable to reject either one out of hand based on these results.

Discussion

For investors who are not capital constrained and who plan to make a number of investments large enough that following on does not significantly affect the number of early stage investments they make, there are



Table 3. A comparison of the performance of all follow-on investments against follow-on investments that happen after a seed round at least doubles.

	All Follow-On Investments with Valuation Changes	Follow-On Investments after Doubling with Valuation Changes
Number	237	45
Fraction with valuation changes	40%	31%
Mean return	2.2x	1.7x
Fraction returning <0.1x ("Went to Zero")	22%	22%
25th percentile return	0.20x	0.14x
50th percentile return	1.25x	1.53x
75th percentile return	2.48x	2.66x

Results shown in this table are gross TVPI values and do not reflect the impact of fees or carry that would reduce returns achievable by an investor.

many reasons outside of returns, strictly construed, that an investor would choose make or avoid follow-on investments. The limiting factor for many investors is the time they spend sourcing and diligencing investments; a follow-on investment offers those investors the opportunity to deploy additional capital without plunging deeply into sourcing and diligencing a new potential investment.

We partly motivated this study based on descriptions of three different world-models, each of which would have suggested different qualitative performance results in simulation. The **Totally Random** world is most supported based on the lack of any clear outcomes and the median/mean split between Never and Always. But the **Momentum** world also appears to have some support: The highest mean return came from Always following on, and it the follow-on investments from seed deals that have doubled seem to be similar to, and typically better than, other follow-on investments, suggesting that performance at conversion is typically a lightly meaningful signal for future performance. Informally, the world model most consistent with our results is something like a mix of 98% **Totally Random** with 2% **Momentum**.



On the other hand, the **Mean-Reverting** world does not seem to have any support from our data, suggesting that *there is no apparent support* for eschewing following on in “hot” Series A deals solely out of a perception that their valuations have been “run up”. At worst these investments should be no different on average than another follow-on investment that happens at a lower multiple, and at best there may be a momentum component to returns that could drive the company’s valuation higher in the future.





Disclosures

This document and the information contained herein is provided for informational and discussion purposes only and is not intended to be a recommendation for any investment or other advice of any kind and shall not constitute or imply any offer to purchase, sell or hold any security or to enter into or engage in any type of transaction. Any such offers will only be made pursuant to formal offering materials containing full details regarding applicable risks, minimum investments, fees, and expenses. Any investment in venture funds, including AngelList funds, involves a high degree of risk and is suitable only for sophisticated and qualified accredited investors.

Past performance is not indicative of future returns. The simulated portfolio returns shown and discussed above do not represent the actual composition and performance of any current or future AngelList fund. These simulated portfolio returns have inherent limitations. They do not reflect the impact that material economic and market factors had or might have had if actual investments had been made. Returns achievable by actual funds may also differ from simulated portfolios due to differences in the timing and prices of investments and the identity and weightings of securities holdings. Data concerning returns from simulated portfolios also do not consider tax impacts, fee waivers, and other factors that would cause actual returns to investors to differ materially from the simulated portfolio returns. While this post presents the simulated performance of several re-investment strategies based on the specified historic AngelList deals, these distributions and probabilities may not be the same during any other period. Some institutional and professional investors on the AngelList platform also have superior access to deals and information from AngelList and leads. The simulated portfolio analysis in this article may include returns from investments in funds that were only available to these institutional and professional investors, which may cause the simulated returns to differ materially from returns achievable by other investors on the AngelList platform. Moreover, market returns based on AngelList data may not be representative of the broader market of early-stage technology startups generally due to a number of factors, including the potential for selection bias in investment opportunities available to AngelList funds and the size and frequency of such investments accessible to a given investor.